

CLAIMS

What is claimed is:

1. A method for member initialization to an asynchronous data replication group in a database system, comprising:

(a) initiating subscription for a new member at an initiator, wherein the initiator is an active member in the replication group;

5 (b) sending a list of active members in the replication group from the initiator to the new member;

(c) sending subscription information for the new member to each active member in the replication group;

(d) validating a subscription to the new member at each active member and
10 sending an acknowledgement to the initiator; and

(e) determining that an acknowledgement from each active member has been received by the initiator.

2. The method of claim 1, further comprising:

15 (f) determining that a loading of data onto the new member is required;

(g) loading the data from the initiator to the new member;

(h) storing current replication messages from the active members onto at least one spill queue;

(i) determining that the loading of the data has completed; and

20 (j) applying the replication messages stored on the at least one spill queue to the new member.

3. The method of claim 1, wherein the initiating (a) comprises:

(a1) inserting a signal into a signal table of the initiator to initiate the subscription for the new member.

5 4. The method of claim 1, wherein the sending (b) comprises:

(b1) computing the list of active members in the replication group by the initiator;

(b2) changing a subscription state for the new member to 'initializing' at the initiator; and

(b3) sending the list of active members from the initiator to the new member.

10 5. The method of claim 1, wherein the sending (c) comprising:

(c1) receiving the list of active members from the initiator by the new member;

(c2) sending the subscription information for the new member from the new member to each active member in the list of active members; and

15 (c3) changing a subscription state at the new member for each active member to 'active'.

6. The method of claim 1, wherein the validating (d) comprises:

(d1) receiving the subscription information for the new member by each active member from the new member;

20 (d2) changing a subscription state for the subscription to the new member to 'active' at each active member, if no loading of data into the new member is required;

(d3) changing the subscription state for the subscription to the new member to

‘loading’ at each active member, if loading of data into the new member is required;

(d4) informing the initiator that changes to the new member is spooling by each active member; and

(d5) informing the new member that its subscription information has been received by each active member.

7. The method of claim 1, wherein the determining (e) comprising:

(e1) collecting messages from the active members by the initiator, wherein each message informs the initiator that the sending active member is spooling changes to the new member; and

(e2) determining that messages from all active members in the replication group have been received by the initiator.

8. The method of claim 2, wherein the determining (f) further comprises:

(f1) creating a spill queue for each active member.

9. The method of claim 2, wherein the determining (i) comprises:

(i1) sending a load done message from the initiator to the new member.

10. The method of claim 2, wherein the applying (j) comprises:

(j1) sending a load done message to each active member by the new member;

(j2) sending an acknowledgement of receipt of the load done message by each active member to the new member; and

(j3) determining that the acknowledgement of receipt has been received from each active member by the new member.

11. The method of claim 1, further comprising:

5 (f) performing asynchronous data replication between the active members in the replication group, wherein the replication group includes the new member.

12. The method of claim 2, further comprising:

10 (k) performing asynchronous data replication between the active members in the replication group, wherein the replication group includes the new member.

13. A method for member deactivation from an asynchronous data replication group in a database system, comprising:

15 (a) initiating deactivation of a member of the replication group at an initiator, wherein the initiator is an active member in the replication group;

(b) sending a stop message from the initiator to the deactivating member;

(c) sending a stop message from the deactivating member to each active member in the replication group;

20 (d) receiving the stop message from the deactivating member by each active member;

(e) stopping a subscription to the deactivating member by each active member;

(f) sending an acknowledgement of receipt of the stop message to the deactivating member by each active member;

(g) receiving the acknowledgement of receipt of the stop message by the deactivating member from an active member and deactivating a corresponding subscription to the sending active member at the deactivating member; and

(h) determining that the acknowledgement of receipt of the stop message from each active member has been received by the deactivating member.

14. The method of claim 13, wherein the initiating (a) comprises:

(a1) initiating deactivation of the deactivating member by inserting a stop signal into a signal table of the initiator.

15. The method of claim 13, wherein the sending (c) further comprises:

(c1) changing a subscription state for each active member at the deactivating member to 'stopping'.

16. A system, comprising:

an initiator node, wherein the initiator node is an active member of an asynchronous data replication group;

a new member node, wherein the new member node is initialized into the replication group by:

initiating a subscription for the new member node at the initiator node,

sending a list of active member nodes in the replication group from the

initiator node to the new member node,

sending subscription information for the new member to each active member

node in the replication group,

validating a subscription to the new member node at each active member node and sending an acknowledgement to the initiator node, and

determining that an acknowledgement from each active member node has
5 been received by the initiator node.

17. A system, comprising:

an initiator node, wherein the initiator node is an active node of an asynchronous data replication group; and

10 a deactivating node, wherein the deactivating node is to be deactivated from the replication group by:

initiating deactivation of the deactivating node at the initiator node,
sending a stop message from the initiator node to the deactivating node,
sending a stop message from the deactivating node to each active node in the

15 replication group,

receiving the stop message from the deactivating node by each active node,
stopping a subscription to the deactivating node by each active node,
sending an acknowledgement of receipt of the stop message to the

deactivating node by each active node,

20 receiving the acknowledgement of receipt of the stop message by the deactivating node from an active node and deactivating a corresponding subscription to the sending active node at the deactivating node, and

determining that the acknowledgement of receipt of the stop message from

each active node has been received by the deactivating node.

18. A computer readable medium with program instructions for member initialization to an asynchronous data replication group in a database system, comprising:

5 (a) initiating subscription for a new member at an initiator, wherein the initiator is an active member in the replication group;

(b) sending a list of active members in the replication group from the initiator to the new member;

10 (c) sending subscription information for the new member to each active member in the replication group;

(d) validating a subscription to the new member at each active member and sending an acknowledgement to the initiator; and

(e) determining that an acknowledgement from each active member has been received by the initiator.

15 19. The medium of claim 18, further comprising:

(f) determining that a loading of data onto the new member is required;

(g) loading the data from the initiator to the new member;

20 (h) storing current replication messages from the active members onto at least one spill queue;

(i) determining that the loading of the data has completed; and

(j) applying the replication messages stored on the at least one spill queue to the new member.

20. The medium of claim 18, wherein the initiating (a) comprises:

(a1) inserting a signal into a signal table of the initiator to initiate the subscription for the new member.

5 21. The medium of claim 18, wherein the sending (b) comprises:

(b1) computing the list of active members in the replication group by the initiator;

(b2) changing a subscription state for the new member to 'initializing' at the initiator; and

(b3) sending the list of active members from the initiator to the new member.

10 22. The medium of claim 18, wherein the sending (c) comprising:

(c1) receiving the list of active members from the initiator by the new member;

(c2) sending the subscription information for the new member from the new member to each active member in the list of active members; and

15 (c3) changing a subscription state at the new member for each active member to 'active'.

23. The medium of claim 18, wherein the validating (d) comprises:

(d1) receiving the subscription information for the new member by each active member from the new member;

20 (d2) changing a subscription state for the subscription to the new member to 'active' at each active member, if no loading of data into the new member is required;

(d3) changing the subscription state for the subscription to the new member to

‘loading’ at each active member, if loading of data into the new member is required;

(d4) informing the initiator that changes to the new member is spooling by each active member; and

(d5) informing the new member that its subscription information has been
5 received by each active member.

24. The medium of claim 18, wherein the determining (e) comprising:

(e1) collecting messages from the active members by the initiator, wherein each message informs the initiator that the sending active member is spooling changes to the new
10 member; and

(e2) determining that messages from all active members in the replication group have been received by the initiator.

25. The medium of claim 19, wherein the determining (f) further comprises:

(f1) creating a spill queue for each active member.
15

26. The medium of claim 19, wherein the determining (i) comprises:

(i1) sending a load done message from the initiator to the new member.

27. The medium of claim 19, wherein the applying (j) comprises:
20

(j1) sending a load done message to each active member by the new member;

(j2) sending an acknowledgement of receipt of the load done message by each active member to the new member; and

(j3) determining that the acknowledgement of receipt has been received from each active member by the new member.

28. The medium of claim 18, further comprising:

5 (f) performing asynchronous data replication between the active members in the replication group, wherein the replication group includes the new member.

29. The medium of claim 19, further comprising:

10 (k) performing asynchronous data replication between the active members in the replication group, wherein the replication group includes the new member.

30. A computer readable medium with program instructions for member deactivation from an asynchronous data replication group in a database system, comprising:

- 15 (a) initiating deactivation of a member of the replication group at an initiator, wherein the initiator is an active member in the replication group;
- (b) sending a stop message from the initiator to the deactivating member;
- (c) sending a stop message from the deactivating member to each active member in the replication group;
- 20 (d) receiving the stop message from the deactivating member by each active member;
- (e) stopping a subscription to the deactivating member by each active member;
- (f) sending an acknowledgement of receipt of the stop message to the deactivating member by each active member;

(g) receiving the acknowledgement of receipt of the stop message by the deactivating member from an active member and deactivating a corresponding subscription to the sending active member at the deactivating member; and

(h) determining that the acknowledgement of receipt of the stop message from
5 each active member has been received by the deactivating member.

31. The medium of claim 30, wherein the initiating (a) comprises:

(a1) initiating deactivation of the deactivating member by inserting a stop signal
into a signal table of the initiator.

10

32. The medium of claim 30, wherein the sending (c) further comprises:

(c1) changing a subscription state for each active member at the deactivating
member to 'stopping'.